Location Systems for Ubiquitous Computing

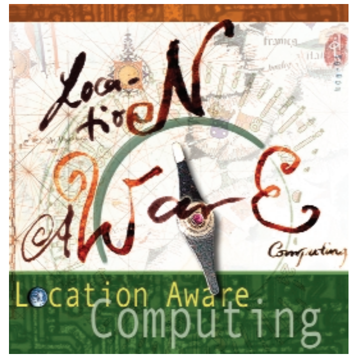

This survey and taxonomy of location systems for mobile-computing applications describes a spectrum of current products and explores the latest research in the field.

Jeffrey Hightower
Gaetano Borriello
University of Washington

To serve us well, emerging mobile computing applications will need to know the physical location of things so that they can record them and report them to us: What lab bench was I standing by when I prepared these tissue samples? How should our search-and-rescue team move to quickly locate all the avalanche victims? Can I automatically display this stock devaluation chart on the large screen I am standing next to?

Researchers are working to meet these and similar needs by developing systems and technologies that automatically locate people, equipment, and other tangibles. Indeed, many systems over the years have addressed the problem of automatic location sensing. Because each approach solves a slightly different problem or supports different applications, they vary in many parameters, such as the physical phenomena used for location determination, the form factor of the sensing apparatus, power requirements, infrastructure versus portable elements, and resolution in time and space.

To make sense of this domain, we have developed a taxonomy to help developers of location-aware applications better evaluate their options when choosing a location-sensing system. The taxonomy may also aid researchers in identifying opportunities for new location-sensing techniques.

LOCATION SYSTEM PROPERTIES

A broad set of issues arises when we discuss and classify location system implementations. These issues are generally independent of the technologies or techniques a system uses, as described in the "Location-Sensing Techniques" sidebar. Although certainly not all orthogonal, nor equally applicable to every system, the classification axes we present do form a reasonable approach to characterizing or evaluating location systems.

The Global Positioning System is perhaps the most widely publicized location-sensing system. GPS provides an excellent lateration framework for determining geographic positions. The worldwide satellite constellation has reliable and ubiquitous coverage and, assuming a differential reference or use of the Wide Area Augmentation System, allows receivers to com-

Location-Sensing Techniques

When attempting to determine a given location, we can choose from three major techniques:

- *Triangulation* can be done via *lateration*, which uses multiple distance measurements between known points, or via *angulation*, which measures angle or bearing relative to points with known separation.
- *Proximity* measures nearness to a known set of points.
- *Scene analysis* examines a view from a particular vantage point.

Location system implementations generally use one or more of these techniques to locate objects, people, or both. A report describing these techniques in detail can be found at www.cs.washington.edu/research/portolano/papers/UW-CSE-01-07-01.pdf.

The resolution of physical positioning systems can have implications for the definitiveness of the symbolic information they can be used to derive.

pute their location to within 1 to 5 meters (http://www.garmin.com/aboutGPS/). Aircraft, hikers, search-and-rescue teams, and rental cars all currently use GPS. Given its celebrity, we use GPS as a running example to introduce our classifiers.

Physical position and symbolic location

A location system can provide two kinds of information: physical and symbolic. GPS provides *physical* positions. For example, our building is situated at 47°39′17″ N by 122°18′23″ W, at a 20.5-meter elevation. In contrast, *symbolic* location encompasses abstract ideas of where something is: in the kitchen, in Kalamazoo, next to a mailbox, on a train approaching Denver.

A system providing a physical position can usually be augmented to provide corresponding symbolic location information with additional information, infrastructure, or both. For example, a laptop equipped with a GPS receiver can access a separate database that contains the positions and geometric service regions of other objects to provide applications with symbolic information.[1] Linking real-time train positions to the reservation and ticketing database can help locate a passenger on a train. Applications can also use the physical position to determine a range of symbolic information. For example, one application can use a single GPS physical position to find the closest printer, while another may link it with calendar information to provide information about that person's current activity.

The distinction between physical position and symbolic location is more pronounced with some technologies than others. GPS is clearly a physical-positioning technology. Point-of-sale logs, bar code scanners, and systems that monitor computer login activity are symbolic location technologies mostly based on proximity to known objects. However, some systems such as Cricket can be used in either mode, depending on their specific configuration.

The resolution of physical-positioning systems can have implications for the definitiveness of the symbolic information they can be used to derive. For example, knowing where a person is inside a building, to within 10 meters, may be ineffective in placing that person in a specific room because of the position of walls within that 10-meter range. Purely symbolic location systems typically provide only very coarse-grained physical positions. Using them often requires multiple readings or sensors to increase accuracy—such as using multiple overlapping proximity sensors to detect someone's position within a room.

Absolute versus relative

An *absolute* location system uses a shared reference grid for all located objects. For example, all GPS receivers use latitude, longitude, and altitude—or their equivalents, such as Universal Transverse Mercator coordinates—for reporting location. Two GPS receivers placed at the same position will report equivalent position readings, and 47°39′17″ N by 122°18′23″ W refers to the same place regardless of GPS receiver.

In a *relative* system, each object can have its own frame of reference. For example, a mountain rescue team searching for avalanche victims can use handheld computers to locate victims' avalanche transceivers. Each rescuer's device reports the victims' position relative to itself.

An absolute location can be transformed into a relative location—relative to a second reference point, that is. However, a second absolute location is not always available. In reverse, we can use triangulation to determine an absolute position from multiple relative readings if we know the absolute position of the reference points. But we often can't know these positions if the reference points are themselves mobile. Thus, the absolute versus relative distinction denotes primarily what information is available and how the system uses it rather than any innate capabilities.

Localized location computation

Some systems provide a location capability and insist that the object being located actually computes its own position. This model ensures privacy by mandating that no other entity may know where the located object is unless the object specifically takes action to publish that information. For example, orbiting GPS satellites have no knowledge about who uses the signals they transmit. Online map servers such as Expedia (http://maps.expedia.com) and old-fashioned road atlases and print maps also fall into this category.

In contrast, some systems require the located object to periodically broadcast, respond with, or otherwise emit telemetry to allow the external infrastructure to locate it. The infrastructure can find objects in its purview without directly involving the objects in the computation. Personal-badge-location systems fit into this category, as do bar codes and the radio frequency identification tags that prevent merchandise theft, track shipments, and help identify livestock in the field (http://www.sensormatic.com and http://www.axsi.com). Placing the burden on the infrastructure decreases the computational and power demands on the objects being located, which makes many more applications possible due to lower costs and smaller form factors.

The policy for manipulating location data need not be dictated by where the computation is performed. For example, system-level access control can provide privacy for a movement history in a personal-location

system while still allowing the infrastructure to perform the location computation. Doing so, however, imposes a requirement of trust in the access control.

Accuracy and precision

A location system should report locations accurately and consistently from measurement to measurement. Some inexpensive GPS receivers can locate positions to within 10 meters for approximately 95 percent of measurements. More expensive differential units usually do much better, reaching 1- to 3-meter accuracies 99 percent of the time. These distances denote the *accuracy*, or grain size, of the position information GPS can provide. The percentages denote *precision*, or how often we can expect to get that accuracy.

Obviously, if we can live with less accuracy, we may be able to trade it for increased precision. Thus, we really must place the two attributes in a common framework for comparison. To arrive at a concise quantitative summary of accuracy and precision, we can assess the error distribution incurred when locating objects, along with any relevant dependencies such as the necessary density of infrastructural elements. For example, "Using five base stations per 300 square meters of indoor floor space, location-sensing system *X* can accurately locate objects within error margins defined by a Gaussian distribution centered at the objects' true locations and having a standard deviation of 2 meters."

Sensor fusion seeks to improve accuracy and precision by integrating many location or positioning systems to form hierarchical and overlapping levels of resolution. Statistically merging error distributions is an effective way to assess the combined effect of multiple sensors.

The ad hoc sensor networking and smart dust community (http://www.darpa.mil/ito/research/sensit) often addresses the related issue of *adaptive fidelity*. A location system with this ability can adjust its precision in response to dynamic situations such as partial failures or directives to conserve battery power.

Often, we evaluate a location-sensing system's accuracy to determine whether it is suitable for a particular application. Motion-capture installations that support computer animation (http://www.ascension-tech.com) feature centimeter-level spatial positioning and precise temporal resolution, but most applications do not require this level of accuracy. GPS tags might suffice for species biologists concerned about the position of a migrating whale pod to a precision of 1 square kilometer. A personal location system for home or office applications might need enough accuracy to answer the query, "Which room was I in around noon?" but not "Where, to the nearest cubic centimeter, was my left thumb at 12:01:34 p.m.?"

Scale

A location-sensing system may be able to locate objects worldwide, within a metropolitan area, throughout a campus, in a particular building, or within a single room. Further, the number of objects the system can locate with a certain amount of infrastructure or over a given time may be limited. For example, GPS can serve an unlimited number of receivers worldwide using 24 satellites plus three redundant backups. On the other hand, some electronic tag readers cannot read *any* tag if more than one is within range.

To assess the scale of a location-sensing system, we consider its coverage area per unit of infrastructure and the number of objects the system can locate per unit of infrastructure per time interval. Time reflects an important consideration because of the limited bandwidth available in sensing objects. For example, a radio-frequency-based technology can only tolerate a maximum number of communications before the channel becomes congested. Beyond this threshold, either latency in determining the objects' positions will increase or a loss in accuracy will occur because the system calculates the objects' positions less frequently.

Systems can often expand to a larger scale by increasing the infrastructure. For example, a tag system that locates objects in a single building can operate on a campus by outfitting all campus buildings and outdoor areas with the necessary sensor infrastructure. Hindrances to scalability in a location system include not only the infrastructure cost but also middleware complexity—it may prove difficult to manage the larger and more distributed databases required for a campus-sized deployment.

For applications that need to recognize or classify located objects to take a specific action based on their location, an automatic identification mechanism is needed.

Recognition

For applications that need to recognize or classify located objects to take a specific action based on their location, an automatic identification mechanism is needed. For example, a modern airport baggage handling system needs to automatically route outbound and inbound luggage to the correct flight or claim carousel. A proximity-location system consisting of tag scanners installed at key locations along the automatic baggage conveyers makes recognition a simple matter of printing the appropriate destination codes on the adhesive luggage check stickers. In contrast, GPS satellites have no inherent mechanism for recognizing individual receivers.

Systems with recognition capability may recognize only some feature types. For example, cameras and vision systems can easily distinguish the color or shape of an object but cannot automatically recognize individual people or a particular apple drawn from a bushel basket.

Figure 1. Olivetti Active Badge (right) and a base station (left) used in the system's infrastructure.

A general technique for providing recognition capability assigns names or globally unique IDs (GUID) to objects the system locates. Once a tag, badge, or label on the object reveals its GUID, the infrastructure can access an external database to look up the name, type, or other semantic information about the object. It can also combine the GUID with other contextual information so it can interpret the same object differently under varying circumstances. For example, a person can retrieve the descriptions of objects in a museum in a specified language. The infrastructure can also reverse the GUID model to emit IDs such as URLs that mobile objects can recognize and use.[2]

Cost

We can assess the cost of a location-sensing system in several ways. Time costs include factors such as the installation process's length and the system's administration needs. Space costs involve the amount of installed infrastructure and the hardware's size and form factor.

Capital costs include factors such as the price per mobile unit or infrastructure element and the salaries of support personnel. For example, GPS receivers need an antenna of sufficient size for adequate satellite reception and may need a second antenna to receive the land-based differential signal. Support personnel at the US Air Force GPS command station must regularly monitor the status of the GPS satellites. Further, building and launching the satellites required a major capital investment by the US government.

A simple civilian GPS receiver costs around $100 and represents the incremental cost of making a new object positionable independently of its global location. A system that uses infrared beacons for broadcasting room IDs requires a beacon for every room in which users want the system to find them. In this case, both the infrastructure and the object the system locates contribute to the incremental cost.

Limitations

Some systems will not function in certain environments. One difficulty with GPS is that receivers usually cannot detect the satellites' transmissions indoors. This limitation has implications for the kind of applications we can build using GPS. For example, because most wired phones are located indoors, even if its accuracy and precision were high enough to make it conceivable, GPS does not provide adequate support for an application that routes phone calls to the land-line phone nearest the intended recipient. A possible solution that maintains GPS interaction yet works indoors uses a system of GPS repeaters mounted at the edges of buildings to rebroadcast the signals inside.

Some tagging systems can read tags properly only when a single tag is present. In some cases, colocated systems that use the same operating frequency experience interference. In general, we assess functional limitations by considering the characteristics of the underlying technologies that implement the location system.

A SURVEY OF LOCATION SYSTEMS

We can use our taxonomy to survey some of the research and commercial location technologies that are representative of the location-sensing field. Table 1 summarizes the properties of these technologies. In the table, the open circles indicate that the systems can be classified as either absolute or relative, and the checkmarks indicate that localized location computation (LLC) or recognition applies to the system. Physical-symbolic and absolute-relative are paired alternatives, and a system is usually one or the other in each category.

Active Badge

The first and arguably archetypal indoor badge sensing system, the Active Badge location system, which was developed at Olivetti Research Laboratory, now AT&T Cambridge,[3] consists of a cellular proximity system that uses diffuse infrared technology. Each person the system can locate wears a small infrared badge like that shown in Figure 1. The badge emits a globally unique identifier every 10 seconds or on demand. A central server collects this data from fixed infrared sensors around the building, aggregates it, and provides an application programming interface for using the data.

The Active Badge system provides absolute location information. A badge's location is symbolic, representing, for example, the room—or other infrared constraining volume—in which the badge is located. The Cambridge group also designed one of the first large software architectures for handling this type of symbolic location data.[4]

As with any diffuse infrared system, Active Badges have difficulty in locations with fluorescent lighting or direct sunlight because of the spurious infrared emissions these light sources generate. Diffuse infrared has an effective range of several meters, which limits cell sizes to small- or medium-sized rooms. In larger rooms, the system can use multiple infrared beacons.

Active Bat

In more recent work, AT&T researchers have developed the Active Bat location system, which uses an ultrasound time-of-flight lateration technique to provide more accurate physical positioning than Active Badges.[5] Users and objects carry Active Bat tags. In

Table 1. Current location sensing technologies.

Technology	Technique	Physical	Symbolic	Absolute	Relative	LLC	Recognition	Accuracy and precision if available	Scale	Cost	Limitations
GPS	Radio time-of-flight lateration	•		•		✓		1-5 meters (95-99 percent)	24 satellites worldwide	Expensive infrastructure $100 receivers	Not indoors
Active Badges	Diffuse infrared cellular proximity		•	•			✓	Room size	1 base per room, badge per base per 10 sec	Administration costs, cheap tags and bases	Sunlight and fluorescent light interfere with infrared
Active Bats	Ultrasound time-of-flight lateration	•		•			✓	9 cm (95 percent)	1 base per 10 square meters, 25 computations per room per sec	Administration costs, cheap tags and sensors	Required ceiling sensor grid
MotionStar	Scene analysis, lateration	•		•			✓	1 mm, 1 ms, 0.1° (nearly 100 percent)	Controller per scene, 108 sensors per scene	Controlled scenes, expensive hardware	Control unit tether, precise installation
VHF Omini-directional Ranging	Angulation	•		•		✓		1° radial (≈ 100 percent)	Several transmitters per metropolitan area	Expensive infrastructure, inexpensive aircraft receivers	30-140 nautical miles, line of sight
Cricket	Proximity, lateration		•	◦	◦	✓		4 × 4 ft. regions (≈ 100 percent)	≈ 1 beacon per 16 square ft.	$10 beacons and receivers	No central management receiver computation
MSR RADAR	802.11 RF scene analysis and triangulation	•		•			✓	3-4.3 m (50 percent)	3 bases per floor	802.11 network installation, ≈ $100 wireless NICs	Wireless NICs required
PinPoint 3D-iD	RF lateration	•		•			✓	1-3 m	Several bases per building	Infrastructure installation, expensive hardware	Proprietary, 802.11 interference
Avalanche Transceivers	Radio signal strength proximity	•			•			Variable, 60-80 meter range	1 transceiver per person	≈ $200 per transceiver	Short radio range, unwanted signal attenuation
Easy Living	Vision, triangulation		•	•			✓	Variable	3 cameras per small room	Processing power, installation cameras	Ubiquitous public cameras
Smart Floor	Physical contact proximity	•		•			✓	Spacing of pressure sensors (100 percent)	Complete sensor grid per floor	Installation of sensor grid, creation of footfall training dataset	Recognition may not scale to large populations
Automatic ID systems	Proximity		•	◦	◦		✓	Range of sensing phenomenon (RFID typically <1m)	Sensor per location	Installation, variable hardware costs	Must know sensor locations
Wireless Andrew	802.11 proximity		•	•			✓	802.11 cell size, (≈ approx. 100 m indoor, 1 km free space)	Many bases per campus	802.11 deployment, ≈ $100 wireless NICs	Wireless NICs required, RF cell geometries
E911	Triangulation	•		•			✓	150-300 m (95 percent)	Density of cellular infrastructure	Upgrading phone hardware or cell infrastructure	Only where cell coverage exists
SpotON	Ad hoc lateration	•			•		✓	Depends on cluster size	Cluster at least 2 tags	$30 per tag, no infrastructure	Attenuation less accurate than time-of-flight

Electromagnetic sensing is the position-tracking technology behind much of the research and many of the products that support virtual reality and motion capture for computer animation.

response to a request the controller sends via short-range radio, a Bat emits an ultrasonic pulse to a grid of ceiling-mounted receivers. At the same time the controller sends the radio frequency request packet, it also sends a synchronized reset signal to the ceiling sensors using a wired serial network. Each ceiling sensor measures the time interval from reset to ultrasonic pulse arrival and computes its distance from the Bat. The local controller then forwards the distance measurements to a central controller, which performs the lateration computation. Statistical pruning eliminates erroneous sensor measurements caused by a ceiling sensor hearing a reflected ultrasound pulse instead of one that traveled along the direct path from the Bat to the sensor.

The system, as reported in 1999, can locate Bats to within 9 cm of their true position for 95 percent of the measurements, and work to improve the accuracy even further is in progress. It can also compute orientation information given predefined knowledge about the placement of Bats on the rigid form of an object and allowing for the ease with which ultrasound is obstructed. Each Bat has a GUID for addressing and recognition.

Using ultrasound time of flight this way requires a large fixed-sensor infrastructure throughout the ceiling and is rather sensitive to the precise placement of these sensors. Thus, scalability, ease of deployment, and cost are disadvantages of this approach.

Cricket

Complementing the Active Bat system,[6] the Cricket Location Support System uses ultrasound emitters to create the infrastructure and embeds receivers in the object being located. This approach forces the objects to perform all their own triangulation computations. Cricket uses the radio frequency signal not only for synchronization of the time measurement, but also to delineate the time region during which the receiver should consider the sounds it receives. The system can identify any ultrasound it hears after the end of the radio frequency packet as a reflection and ignore it. A randomized algorithm allows multiple uncoordinated beacons to coexist in the same space. Each beacon also transmits a string of data that describes the semantics of the areas it delineates using the short-range radio.

Like the Active Bat system, Cricket uses ultrasonic time-of-flight data and a radio frequency control signal, but this system does not require a grid of ceiling sensors with fixed locations because its mobile receivers perform the timing and computation functions. Cricket, in its currently implemented form, is much less precise than Active Bat in that it can accurately delineate 4 × 4 square-foot regions within a room, while Active Bat is accurate to 9 cm. However, the fundamental limit of range-estimation accuracy used in Cricket should be no different than Active Bat, and future implementations may compete with each other on accuracy.

Cricket implements both the lateration and proximity techniques. Receiving multiple beacons lets receivers triangulate their position. Receiving only one beacon still provides useful proximity information when combined with the semantic string the beacon transmits on the radio.

Cricket's advantages include privacy and decentralized scalability, while its disadvantages include a lack of centralized management or monitoring and the computational burden—and consequently power burden—that timing and processing both the ultrasound pulses and RF data place on the mobile receivers.

RADAR

A Microsoft Research group has developed RADAR,[7] a building-wide tracking system based on the IEEE 802.11 WaveLAN wireless networking technology.

RADAR measures, at the base station, the signal strength and signal-to-noise ratio of signals that wireless devices send, then it uses this data to compute the 2D position within a building. Microsoft has developed two RADAR implementations: one using scene analysis and the other using lateration.

The RADAR approach offers two advantages: It requires only a few base stations, and it uses the same infrastructure that provides the building's general-purpose wireless networking. Likewise, RADAR suffers two disadvantages. First, the object it is tracking must support a wireless LAN, which may be impractical on small or power-constrained devices. Second, generalizing RADAR to multifloored buildings or three dimensions presents a nontrivial problem.

RADAR's scene-analysis implementation can place objects to within about 3 meters of their actual position with 50 percent probability, while the signal-strength lateration implementation has 4.3-meter accuracy at the same probability level. Although the scene-analysis version provides greater accuracy, significant changes in the environment, such as moving metal file cabinets or large groups of people congregating in rooms or hallways, may necessitate reconstructing the predefined signal-strength database or creating an entirely new database.

Several commercial companies such as WhereNet (http://www.widata.com) and Pinpoint (http://www.pinpointco.com) sell wireless asset-tracking packages, which are similar in form to RADAR. Pinpoint's 3D-iD performs indoor position tracking using proprietary base station and tag hardware to measure radio time of flight. Pinpoint's system achieves 1- to 3-meter

Figure 2. MotionStar DC magnetic tracker, a precision system used in motion capture for computer animation, tracks the position and orientation of up to 108 sensor points on an object or scene. Key components include (left and right) the magnetic pulse transmitting antennas and (center) the receiving antennas and controller. Image courtesy of Ascension Technology Corporation.

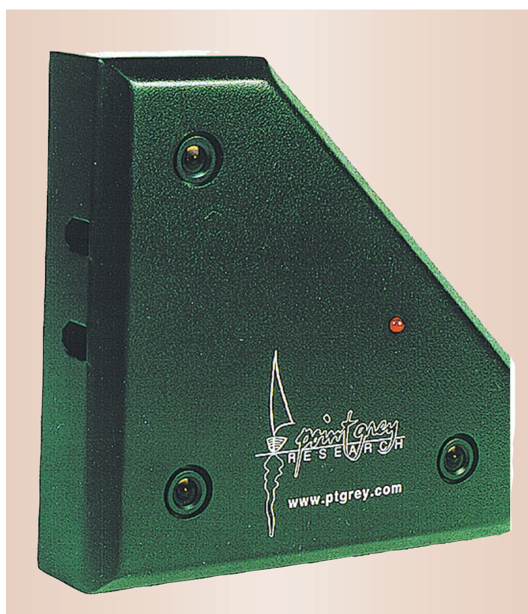

Figure 3. Digiclops color 3D camera, made by Point Grey Research and used by the Microsoft Research Easy Living group to provide stereo-vision positioning in a home environment. Image courtesy of Point Grey Research Inc.

accuracy and, by virtue of being a commercial product, offers easier deployment and administration than many research systems.

The 3D-iD system suffers the disadvantage that each antenna has a narrow cone of influence, which can make ubiquitous deployment prohibitively expensive. Thus, 3D-iD best suits large indoor space settings such as hospitals or warehouses. It has difficulty interoperating with the 802.11 wireless networking infrastructure because of radio spectrum collision in the unregulated Industrial, Scientific, and Medical band.

MotionStar magnetic tracker

Electromagnetic sensing offers a classic position-tracking method.[8] The large body of research and products that support virtual reality and motion capture for computer animation often offer modern incarnations of this technology. For example, Ascension offers a variety of motion-capture solutions, including Flock of Birds and, shown in Figure 2, the MotionStar DC magnetic tracker.[9] These tracking systems generate axial DC magnetic-field pulses from a transmitting antenna in a fixed location. The system computes the position and orientation of the receiving antennas by measuring the response in three orthogonal axes to the transmitted field pulse, combined with the constant effect of the earth's magnetic field.

Tracking systems such as MotionStar sense precise physical positions relative to the magnetic transmitting antenna. These systems offer the advantage of very high precision and accuracy, on the order of less than 1 mm spatial resolution, 1 ms time resolution, and 0.1° orientation capability. Disadvantages include steep implementation costs and the need to tether the tracked object to a control unit. Further, the sensors must remain within 1 to 3 meters of the transmitter, and accuracy degrades with the presence of metallic objects in the environment.

Many other technologies have been used in virtual environments or in support of computer animation. A CDMA radio ranging approach has been suggested,[10] and many companies sell optical, infrared, and mechanical motion-capture systems. Like MotionStar, these systems are not designed to be scalable for use in large, location-aware applications. Rather, they capture position in one precisely controlled environment.

Easy Living

Several groups have explored using computer vision technology to figure out where things are. Microsoft Research's Easy Living provides one example of this approach. Easy Living uses the Digiclops real-time 3D cameras shown in Figure 3 to provide stereo-vision positioning capability in a home environment.[11] Although Easy Living uses high-performance cameras, vision systems typically use substantial amounts of processing power to analyze frames captured with comparatively low-complexity hardware.

State-of-the-art integrated systems[12] demonstrate that multimodal processing—silhouette, skin color, and face pattern—can significantly enhance accuracy. Vision location systems must, however, constantly struggle to maintain analysis accuracy as scene complexity increases and more occlusive motion occurs.

Figure 4. Robots have many onboard sensors for use in localization, multirobot collaboration, and zero-knowledge map building.

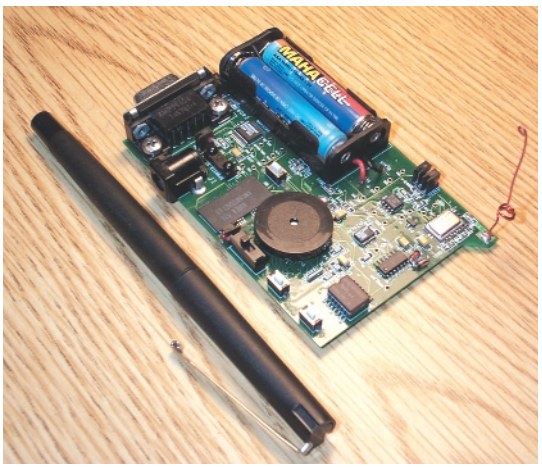

Figure 5. Prototype SpotON radio tag. These tags use radio signal attenuation to perform ad hoc lateration. Ad hoc clusters of tags cooperate to factor out measurement errors for all tag positions.

The dependence on infrastructural processing power, along with public wariness of ubiquitous cameras, can limit the scalability or suitability of vision location systems in many applications.

Smart Floor

In Georgia Tech's Smart Floor proximity location system, embedded pressure sensors[13] capture footfalls, and the system uses the data for position tracking and pedestrian recognition. This unobtrusive direct physical contact system does not require people to carry a device or wear a tag. However, the system has the disadvantages of poor scalability and high incremental cost because the floor of each building in which Smart Floor is deployed must be physically altered to install the pressure sensor grids.

E911

The US Federal Communications Commission's E911 telecommunication initiatives require that wireless phone providers develop a way to locate any phone that makes a 911 emergency call (http://www.fcc.gov/e911). E911 is not a specific location-sensing system, but we include it because the initiatives have spawned many companies that are developing a variety of location systems to determine a cellular phone's location.

Location systems developed to comply with the E911 initiatives will also support new consumer services. For example, a wireless telephone can use this technology to find the nearest gas station, post office, movie theater, bus, or automated teller machine. Data from many cellular users can be aggregated to identify areas of traffic congestion. Many business speculators tout this model of mobile consumerism, or *m-commerce*, as being the "next big thing."

To comply with E911, vendors are exploring several RF techniques, including antenna proximity, angulation using phased antenna arrays, lateration via signal attenuation and time of flight, as well as GPS-enabled handsets that transmit their computed location to the cellular system (http://www.airbiquity.com). To meet the FCC requirement, positioning must be accurate to within 150 meters for 95 percent of calls with receiver-based handset solutions such as GPS, or to within 300 meters with network-transmitter-based approaches.

RESEARCH DIRECTIONS

Location sensing is a mature enough field to define a space within a taxonomy that is generally populated by existing systems, as Table 1 shows. As such, future work should generally focus on lowering cost, reducing the amount of infrastructure, improving scalability, and creating systems that are more flexible within the taxonomy. This does not imply, however, that location sensing is a solved problem or that further advancements are simply a matter of rote technology improvement. Rather, location sensing is now entering an exciting phase in which cross-pollination with ideas from other computer science and engineering disciplines motivates future research.

Sensor fusion

Defined as the use of multiple technologies or location systems simultaneously to form hierarchical and overlapping levels of sensing, sensor fusion can provide aggregate properties unavailable when using location systems individually.

For example, integrating several systems with different error distributions may increase accuracy and precision beyond what is possible using an individual system. The more independent the techniques, the more effectively they can be combined.

An example of current sensor fusion research, multisensor collaborative robot localization and map building presents a problem usually divided into two subproblems:

- tracking location as the environment changes or the robot moves, and
- determining robot location from a zero-knowledge start state.

Autonomous robots, such as those shown in Figure 4, employ a myriad of onboard sensors including ultrasound and laser range finders, inertial trackers, and cameras. The robots use Markov and Bayesian statistical techniques and multirobot collaboration to accomplish sensor fusion.[14] These techniques provide important starting points for combining location systems for ubiquitous computing.

Ad hoc location sensing

This approach to locating objects without drawing on the infrastructure or central control borrows ideas from the ad hoc networking research community. In a purely ad hoc location-sensing system, all of the entities become mobile objects with the same sensors and capabilities. To estimate their locations, objects cooperate with other nearby objects by sharing sensor data to factor out overall measurement error. In this way, a cluster of ad hoc objects converges to an accurate estimate of all nearby objects' positions. Objects in the cluster are located relative to one another or absolutely if some objects in the cluster occupy known locations.

The techniques for building ad hoc systems include triangulation, scene analysis, or proximity. The work of Lance Doherty and colleagues[15] and Nirupama Bulusu and colleagues[16] explores ad hoc proximity systems that consider variants of the following question: Given a set S of tiny sensor devices and a proximity model of radio connectivity, such as a sphere or circle with a fixed radius, if we know that $s_0 \ldots s_n, s_i \subset S$ are subsets of sensors in proximity to one another, how accurately can we infer the relative location of all sensors in set S?

Doherty and colleagues present an algorithmic approach to this problem as well as a framework for describing error bounds on the computed locations. Bulusu and colleagues extend this basic connectivity notion by adding an ideal theoretical model of outdoor radio behavior and a regular grid of reference nodes at known locations.

The SpotON system implements ad hoc lateration with low-cost tags. SpotON tags use radio signal attenuation to estimate intertag distance.[17] They exploit the density of tags and correlation of multiple measurements to improve both accuracy and precision. Figure 5 shows a prototype SpotON tag.

Sensing object locations with no fixed infrastructure represents a highly scalable and low-cost approach. In the future, infrastructural systems could incorporate ad hoc concepts to increase accuracy or reduce cost. For example, it might be possible for a system like Active Bat to use a sparser ceiling-mounted ultrasound receiver grid if Bats could also accurately measure their distance from other Bats and share this information with the infrastructure.

Location-sensing-system accuracy: A challenge

Comparing the accuracy and precision of different location sensing systems can be an arduous task because many system descriptions lack a concise summary of these parameters. We therefore suggest that future quantitative evaluations of location-sensing systems include the *error distribution*, summarizing the system's accuracy and precision and any relevant dependencies such as the density of infrastructural elements. For example, "Using five base stations per 300 square meters of indoor floor space, location-sensing system *X* can accurately locate objects within error margins defined by a Gaussian distribution centered at the objects' true location and a standard deviation of 2 meters." We strongly encourage the location-sensing research and development community to investigate how to best obtain and represent such error distributions.

In addition to its comparison value, researchers could use a location-sensing system's accurately described error distribution as partial input for simulating a system—even a hypothetical one. Prototyping an application with a simulator avoids the cost of purchasing, deploying, and configuring a hardware infrastructure when the goal is simply to evaluate the suitability of a certain location-sensing system. Preliminary work on this idea has begun. For example, Markus Bylund and Fredrik Espinoza have built a simulator for a campus-sized position-sensing system that uses a Quake III gaming arena.[18]

Researchers can apply our taxonomy to evaluate the characteristics of the location system a particular application needs, or they can use it to help determine the suitability of an existing location system for that application. With decreasing costs of silicon and wireless connectivity, location systems will become increasingly common. Increased attention and effort will foster improvements in various aspects of the design space. We offer our approach to comparing these systems to help researchers make better choices for the location systems they use in ubiquitous applications. ✷

Acknowledgments

The authors thank Trevor Pering from Intel Research and Ken Yasuhara, Neil Spring, and Vibha Sazawal from the University of Washington for their editorial feedback on this article. We also thank Dieter Fox and Larry Arnstein at UW for providing valuable insights that helped clarify our presentation.

References

1. B. Brumitt et al., "Ubiquitous Computing and the Role of Geometry," *IEEE Personal Comm.*, Oct. 2000, pp. 41-43.

2. J. Barton and T. Kindberg, *The CoolTown User Experience*, tech. report 2001-22, HP Laboratories, Palo Alto, Calif., 2001.
3. R. Want et al., "The Active Badge Location System," *ACM Trans. Information Systems*, Jan. 1992, pp. 91-102.
4. A. Harter and A. Hopper, "A Distributed Location System for the Active Office," *IEEE Network*, Jan./Feb. 1994, pp. 62-70.
5. A. Harter et al., "The Anatomy of a Context-Aware Application," *Proc. 5th Ann. Int'l Conf. Mobile Computing and Networking* (Mobicom 99), ACM Press, New York, 1999, pp. 59-68.
6. N.B. Priyantha, A. Chakraborty, and H. Balakrishnan, "The Cricket Location-Support System," *Proc. 6th Ann. Int'l Conf. Mobile Computing and Networking* (Mobicom 00), ACM Press, New York, 2000, pp. 32-43.
7. P. Bahl and V. Padmanabhan, "RADAR: An In-Building RF-Based User Location and Tracking System," *Proc. IEEE Infocom 2000*, IEEE CS Press, Los Alamitos, Calif., 2000, pp. 775-784.
8. F. Raab et al., "Magnetic Position and Orientation Tracking System," *IEEE Trans. Aerospace and Electronic Systems*, Sept. 1979, pp. 709-717.
9. *Technical Description of DC Magnetic Trackers*, Ascension Technology Corp., Burlington, Vt., 2001.
10. S.R. Bible, M. Zyda, and D. Brutzman, "Using Spread-Spectrum Ranging Techniques for Position Tracking in a Virtual Environment," *Second IEEE Workshop Networked Realities*, http://www.npsnet.org/~zyda/pubs/NR95-Paper-Bible.pdf.
11. J. Krumm et al., "Multi-Camera Multi-Person Tracking for Easy Living," *Proc. 3rd IEEE Int'l Workshop Visual Surveillance*, IEEE Press, Piscataway, N.J., 2000, pp. 3-10.
12. T. Darrell et al., "Integrated Person Tracking Using Stereo, Color, and Pattern Detection," *Conf. Computer Vision and Pattern Recognition*, IEEE CS Press, Los Alamitos, Calif., 1998, pp. 601-608.
13. R.J. Orr and G.D. Abowd, "The Smart Floor: A Mechanism for Natural User Identification and Tracking," *Proc. 2000 Conf. Human Factors in Computing Systems* (CHI 2000), ACM Press, New York, 2000.
14. D. Fox et al., "A Probabilistic Approach to Collaborative Multi-Robot Localization," *Autonomous Robots*, June 2000, pp. 325-244.
15. L. Doherty et al., "Convex Position Estimation in Wireless Sensor Networks," *Proc. Infocom 2001*, IEEE CS Press, Los Alamitos, Calif., 2001.
16. N. Bulusu, J. Heidemann, and D. Estrin, "GPS-Less Low-Cost Outdoor Localization for Very Small Devices," Special Issue on Smart Spaces and Environments, *IEEE Personal Comm.*, Oct. 2000, pp. 28-34.
17. J. Hightower, R. Want, and G. Borriello, *SpotON: An Indoor 3d Location Sensing Technology Based on RF Signal Strength*, UW CSE 2000-02-02, Univ. Washington, Seattle, Feb. 2000.
18. M. Bylund and F. Espinoza, "Using Quake III Arena to Simulate Sensors and Actuators When Evaluating and Testing Mobile Services," *CHI 2001 Extended Abstracts*, ACM Press, New York, Mar./Apr. 2001, pp. 241-242.

Jeffrey Hightower *is a PhD candidate at the University of Washington. His research interests include novel uses of wireless technology such as ad hoc location sensing, software engineering for ubiquitous computing, and designing the user experience of embedded systems. He received an MS in computer science and engineering from the University of Washington. Contact him at jeffro@cs.washington.edu.*

Gaetano Borriello *is a professor at the University of Washington and director of the Intel research lab in Seattle. His research interests include design, development, and deployment of computing systems with particular emphasis on mobile and ubiquitous devices and their application in making life simpler by being as invisible as possible to their owners, being highly specialized and thus highly efficient for the task at hand, and able to exploit their connections to each other and the greater worldwide networks. He received a PhD in computer science from the University of California at Berkeley. Contact him at gaetano@cs.washington.edu.*